

4

LABORATORY FOR  
COMPUTER SCIENCE




MASSACHUSETTS  
INSTITUTE OF  
TECHNOLOGY

AD-A200 988

MIT/LCS/TR-420

Polychannel Systems for  
Mass Digital Communication

David K. Gifford



545 TECHNOLOGY SQUARE, CAMBRIDGE, MASSACHUSETTS 02139

This document has been approved  
for public release and does not  
contain any classification or control markings.

MIT/LCS/TR-420

**Polychannel Systems for  
Mass Digital Communication**

David K. Gifford

July 1988

DTIC  
ELECTE  
S DEC 07 1988 D  
E

This research was supported by the Defense Advanced Research Projects Agency of the Department of Defense and was monitored by the Office of Naval Research under contract number N00014-83-K-0125.

© 1988 Massachusetts Institute of Technology

This document has been approved  
for public release and retains its  
classification if indicated.

88 12 6 088

## REPORT DOCUMENTATION PAGE

1a. REPORT SECURITY CLASSIFICATION Unclassified		1b. RESTRICTIVE MARKINGS													
2a. SECURITY CLASSIFICATION AUTHORITY		3. DISTRIBUTION/AVAILABILITY OF REPORT Approved for public release; distribution is unlimited.													
2b. DECLASSIFICATION/DOWNGRADING SCHEDULE															
4. PERFORMING ORGANIZATION REPORT NUMBER(S) MIT/LCS/TR-420		5. MONITORING ORGANIZATION REPORT NUMBER(S) N00014-83-K-0125													
6a. NAME OF PERFORMING ORGANIZATION MIT Laboratory for Computer Science	6b. OFFICE SYMBOL (If applicable)	7a. NAME OF MONITORING ORGANIZATION Office of Naval Research/Department of Navy													
6c. ADDRESS (City, State, and ZIP Code) 545 Technology Square Cambridge, MA 02139		7b. ADDRESS (City, State, and ZIP Code) Information Systems Program Arlington, VA 22217													
8a. NAME OF FUNDING/SPONSORING ORGANIZATION DARPA/DOD	8b. OFFICE SYMBOL (If applicable)	9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER													
8c. ADDRESS (City, State, and ZIP Code) 1400 Wilson Blvd. Arlington, VA 22217		10. SOURCE OF FUNDING NUMBERS PROGRAM ELEMENT NO.      PROJECT NO.      TASK NO.      WORK UNIT ACCESSION NO.													
11. TITLE (Include Security Classification) <u>POLYCHANNEL SYSTEMS FOR MASS DIGITAL COMMUNICATION</u>															
12. PERSONAL AUTHOR(S) <u>Gifford, David K.</u>															
13a. TYPE OF REPORT Technical	13b. TIME COVERED FROM _____ TO _____	14. DATE OF REPORT (Year, Month, Day) 1988 July	15. PAGE COUNT 26												
16. SUPPLEMENTARY NOTATION															
17. COSATI CODES <table border="1"><tr><th>FIELD</th><th>GROUP</th><th>SUB-GROUP</th></tr><tr><td></td><td></td><td></td></tr><tr><td></td><td></td><td></td></tr><tr><td></td><td></td><td></td></tr></table>	FIELD	GROUP	SUB-GROUP										18. SUBJECT TERMS (Continue on reverse if necessary and identify by block number) Design, Experimentation, Human Factors polychannel system, digital broadcast, content label, query routing, update routing, active medium		
FIELD	GROUP	SUB-GROUP													
19. ABSTRACT (Continue on reverse if necessary and identify by block number) <p>We describe a new type of distributed computer system that looks beyond workstation and local area network assumptions towards the time when computers will be used by everyone at both home and the office. This new system is designed to provide sophisticated information services to an entire metropolitan area. We have combined digital broadcast channels and duplex communication channels in a polychannel system that uses predicate based database content labels to automatically route queries. Our thesis is that a polychannel system produces substantial cost and system scaling advantages while retaining the flexibility of client-server style duplex communication. Experimental data from a two year test of the Boston Community Information System system with hundreds of users supports this thesis.</p>				20. DISTRIBUTION/AVAILABILITY OF ABSTRACT <input checked="" type="checkbox"/> UNCLASSIFIED/UNLIMITED <input type="checkbox"/> SAME AS RPT. <input type="checkbox"/> DTIC USERS	21. ABSTRACT SECURITY CLASSIFICATION Unclassified										
22a. NAME OF RESPONSIBLE INDIVIDUAL <u>Judy Little, Publications Coordinator</u>	22b. TELEPHONE (Include Area Code) <u>(617) 253-5894</u>	22c. OFFICE SYMBOL													

DD FORM 1473, 24 MAR

83 APR edition may be used until exhausted.

All other editions are obsolete.

SECURITY CLASSIFICATION OF THIS PAGE

U.S. Government Printing Office 1985-507-047

Unclassified

## Abstract

We describe a new type of distributed computer system that looks beyond workstation and local area network assumptions towards the time when computers will be used by everyone at both home and the office. This new system is designed to provide sophisticated information services to an entire metropolitan area. We have combined digital broadcast channels and duplex communication channels in a *polychannel* system that uses predicate based database *content labels* to automatically route queries. Our thesis is that a polychannel system produces substantial cost and system scaling advantages while retaining the flexibility of client-server style duplex communication. Experimental data from a two year test of the Boston Community Information System system with hundreds of users supports this thesis.

**Categories and Subject Descriptions:** C.0 [Computer Systems]–General: *System architectures*; C.2.1 [Computer Communication Networks]– Distributed Systems: *Distributed Databases*; H.3.4 [Information Storage and Retrieval]– Systems and Software: *Information Networks*; K.8 [Personal Computing]

**General Terms:** Design, Experimentation, Human Factors

**Additional Key Words and Phrases:** polychannel system, digital broadcast, content label, query routing, update routing, active medium

## 1 Polychannel Systems Provide Important Advantages

The overall goal of the Boston Community Information System project is to explore a new type of computer system architecture capable of providing sophisticated information services to every home and office in a city of perhaps a million people. As part of the project we have designed and implemented a system that addresses this goal, and tested it on a sample population in the Boston area.

We are interested in metropolitan area information systems because we believe that in the future computers will be the basis of a new mass communication medium far more powerful and important than the telephone. Computers have the intrinsic ability to autonomously react to incoming information, which permits computers to be an *active medium* for information as opposed to a *passive medium*, such as the telephone. For example, computers can inform us of important external events or take direct action based upon such events, according to our preprogrammed instructions. Furthermore, computers can be used to interpret documents that include computationally active elements.

In order to explore the potential of computers as a communication medium we started the Boston CommInS project in September, 1982. The central idea of the project was to use both digital broadcast for mass data communication and duplex communication in a single unified framework. Our thesis was that this single framework would integrate the cost and scale advantages of digital broadcast with the flexibility of duplex communication.

Our project has resulted in a new approach to distributed computer systems which we believe will substantially broaden their application. Existing distributed computer systems provide each computer with a duplex network attachment used to communicate with other computers. We propose to add to this basic architecture one or more low-cost digital broadcast facilities that will provide mass digital communication to all of the computers in a distributed system. Our thesis is that the provision of a digital broadcast channel will permit distributed systems to scale in size beyond the limits of present technology and to offer new kinds of services.

We call systems that include both simplex and duplex channels *polychannel* systems. Each remote site in a polychannel system has a local database

used to process most user requests. At each remote site this local database is maintained by a computer that selectively applies updates received from a digital broadcast channel. A central site sends database updates of general interest over the digital broadcast channel to all sites. Whenever a site's database cannot process a portion of a request, duplex communication is used to contact one or more remote server computers.

In general, when a large number of sites are interested in similar data a system designer can use a polychannel system instead of a traditional distributed system to reduce communication and server capacity requirements. Application areas include electronic mail systems, nationwide name servers, software distribution, and distributed database systems.

For example, in order to implement a polychannel airline reservation system we could provide each travel agency with a computer that would maintain a database of unsold seats. A travel agency's computer would keep its local copy of the unsold seat database up-to-date by applying updates received via a broadcast channel. Each travel agency's computer would only communicate with the central airline computer when it actually wanted to book or change a reservation.

An important benefit of a polychannel airline system is that a travel agency can implement sophisticated applications that require high-bandwidth access to the database of unsold seats. Yet there is no need for the airline to provide additional server or communication capacity to accommodate an agency's applications. Furthermore the load presented to the airline system is reduced from that of the standard duplex system approach.

In the remainder of this paper we define a polychannel system and compare it with previous work (Section 2), describe our prototype polychannel system, the Boston Community Information System (Section 3), report experimental data from a test of the Boston Community Information System (Section 4), and conclude with our view of how polychannel systems may change the way we view shared information (Section 5).

## 2 The Polychannel Architecture

Previous work that is related to polychannel systems can be divided into two major categories:

- *Simplex* systems use digital broadcast transmission to send information to terminals or computers. Because simplex systems are inherently unidirectional they are limited to information distribution applications. However, they scale without limit since an incremental user site places no additional load on the distribution system. This zero incremental cost also provides simplex systems with a natural economic advantage.

Simplex systems implemented before 1982 were either Teletex systems or stock-quote services. Teletex [Tanton79] sends a set of preformatted pages to terminals via broadcast transmissions. Stock-quote services such as Lotus' Signal [Lotus87] transmit financial data over simplex channel to personal computers where the data is processed and presented. We implemented the first text-based simplex system that used personal computers in 1984. In the last few years additional systems similar to ours have become operational, including X\*Press, Mainstream Data, and Stargate. X\*Press Information Services transmits news information via cable television to personal computers, Mainstream Data provides a general purpose digital broadcast service to major cities in the United States, and Stargate uses the satellite TV station WTBS to transmit electronic bulletin boards to remote computer sites.

- *Duplex* systems offer interactive access to server computers via circuit switched connections (such as the telephone network) or packet-switched networks. Duplex systems can provide a full range of services but server capacity must grow linearly with the number of system users. Thus duplex systems do not have the same economy of scale found in simplex systems.

All commercial information retrieval systems use duplex modem based communication to connect terminals to a central time-sharing system. For example, France's Minitel, Dow-Jones' News Retrieval, and Mead Data Central's NEXIS all work on this principle.

Our use of duplex systems builds on recent work in distributed database systems. In a distributed database system (such as System R\* [Lindsay84]) duplex communication is used to enable data distribution among a confederation of cooperating computers.

A polychannel system is designed to integrate the advantages of duplex and simplex systems. A polychannel system is comprised of:

- *Personal computers*, one per user, for prescreening, processing, and presenting information.
- *Server computers* that implement services such as historical databases, transactional services, and electronic mail. The servers can be operated by many different organizations, and can be decentralized.
- One or more *mass digital communication channels* that broadcast information to the personal computers. Information can be transmitted via cable TV, direct broadcast satellite, subcarriers on normal TV or FM stations, baseband RF broadcast, local area networks, or on-hook ISDN subscriber loops. The information broadcast can include complete data records and programs for integration into a personal computer's local database, indices and abstracts of services available on server computers, electronic mail notifications, and server status information.
- A *duplex communication system* that permits computers to access each other on demand. The duplex communication system can be circuit or packet switched.
- System software that provides *location transparency*. A system is location transparent when it hides from a user where and how services are implemented. When a system is location transparent it can freely distribute data and processing in order to minimize the cost of request processing and to maximize the performance of the system. For example, we expect that in a polychannel system most user transactions will be performed by the user's personal computer. Users may be unaware that certain requests will be partially or completely processed by servers. The next section describes in detail how the Boston Community Information System implements location transparency.

The advantages that polychannel systems offer over duplex or simplex systems can be summarized as follows:

- Polychannel systems scale beyond the limits of duplex systems. In a duplex system server capacity must grow linearly with the number of

users because each user request is processed by a central server. Our experimental data suggest that over .99 of user information retrieval transactions can be locally processed in a polychannel system. Thus polychannel systems will require substantially less centralized server capacity than a duplex system.

- Polychannel systems offer significant cost advantages over duplex systems. Digital broadcast has a zero incremental cost per additional user. Thus information sent via broadcast can be made available at a relatively low price.
- Polychannel systems permit users to add substantial value to broadcast data using preexisting personal computers. This potentially further lowers the cost of a polychannel system.
- Polychannel systems provide a much richer set of services than do simplex systems. In simplex systems there is no mechanism for users to communicate with one another or with a central server. Thus many important applications such as electronic mail and transaction processing that can not be provided in the framework of a simplex system.
- Polychannel systems provide a higher degree of privacy than do duplex systems. Because a polychannel system implements information filtering in a user's personal computer, a user's interest profile is never transmitted to a central server. Thus it is impossible for an intruder to gain access to this interest profile without physical access to the user's personal computer.

### **3 The Boston Community Information System**

The requirements of a community information system suggest a polychannel based architecture and well match the advantages of polychannel systems: a typical community based system will need to support over  $10^5$  users; users will want an inexpensive system; the system will need to support a broad spectrum of information and transaction based applications; and users will want to ensure that their detailed interests remain totally private.




Figure 1: Boston Community Information System Block Diagram

We built the *Boston Community Information System* (BCIS) as an experimental test of our thesis that polychannel systems can implement a personal dynamic medium for an entire metropolitan area. The BCIS provides up-to-the-minute access to the New York Times and Associated Press. In addition to the New York Times and Associated Press information we ran limited tests that provided electronic mail, USENET bulletin boards, and programs via BCIS.

Figure 1 is a schematic diagram of the BCIS system.

- *Server computers* at our laboratory transmit database updates to remote sites and service user transactions. These servers continuously receive information from the New York Times and the Associated Press. As this information is received it is queued for transmission via our dig-

ital broadcast channel and is stored in a full-text database system. We receive over 3 MBytes of data per day from our information suppliers. High server availability and reliability are achieved by our use of two server computers in an alternate-main configuration.

- We broadcast articles via a *digital radio channel* as soon as the articles are received by our servers, and then periodically thereafter for 2 days. The channel carries 20 MBytes of articles per day. The broadcast channel we use is the 92 KHz FM SCA channel on WERS-FM, which can be received within 10 to 20 miles from downtown Boston. Over this SCA channel we transmit asynchronous serial data (8 data bits, 1 start bit, 1 stop bit) at 4.8 KBits/second with frequency shift keying (0 is transmitted as 89 KHz and a 1 is transmitted as 95 KHz) [Gifford85a]. The packet radio receiver for this channel has an RS-232 output, is about the size of a small paperback book, and costs less than \$100 to manufacture (including all cabling).
- Each user has a *personal computer* that responds to requests and maintains a local database drawn from the digital broadcasts. Requests are forwarded to a server computer when they can not be processed locally. The contents of the local database is defined by the user via a user specified *filter* that describes what information should be locally retained. A filter thus provides each user with the ability to define what the front page of his or her own personal newspaper will contain.
- Certain personal computers also include auto-dial *modems* that permit the personal computers to forward queries to a server for processing. A user will notice a delay when a query is forwarded to a server, but otherwise the use of server computers is completely transparent.

In the following sections we describe the system's user interface, how content labels are used to route queries, and our communication protocols. Further details on the architecture of the system can be found in [Gifford85b].

#### *The User Interface*

The user interface to the BCIS System is shown in Figures 2, 3, and 4. The filter, summary, and article windows form the primary interface to the

```

Filter Window (line 4) Viewing lines 4 through 22 of 22     aug 20, 12:40
( 1) (type: nyt) (subject: newssummary)
( 1) (subject: briefing)
( 1) (subject: findingest)
( 1) (subject: business digest)
( 1) (subject: wall street)
( 1) (subject: patents or technology)
( 1) (subject: telerate or rates)
( 1) (subject: forecast# boston vicinity)
( 1) (subject: calendar)
( 1) (subject: movie#)
((date: (date -90 : 0)) and (type: nyt))
( 1) ibm or compaq or (apple and com Cancel Query (ESC)
( 1) mits or "data general" or "digi
( 1) "star wars" or sdi or "strategic defense"
( 1) health@mit
( 1) mit or "massachusetts institute"
( 1) harvard (category: not sports)
( 1) aids
( 1) (category: news) (priority: urgent | bulletin | flash)
( 1) #

```

TERMINAL: The receiver is running on port 1  
 Welcome to the MIT Community Information System - Type F1 for help

Figure 2: Filter Window

system, and show successively more detail about topics. A user can switch between these windows with a single keystroke. Pop-up menus and keyboard commands are used in a uniform way in all of the windows. For example, a carriage-return processes the current query regardless of what window is in view.

The filter window (Figure 2) contains a prioritized listing of a user's interests. The user edits the filter window to determine the information that his or her personal computer will retain from the digital broadcast stream. The filter window serves as the "front page" of a user's personalized newspaper and shows how many articles have been received by filter line. Filter lines can easily be edited and reprioritized with keyboard commands. A user can set triggers on certain filter lines so that when a preselected article arrives the computer will either notify the user or it will print out the article.

The summary window (Figure 3) contains a short summary of each article satisfying the current query. This summary is composed of the first sentence of each matching article.

The article window (Figure 4) displays the currently selected article. This window allows articles to be viewed, printed, and deleted.

The performance of the system is quite good. On an IBM AT sub-second response is provided for local query processing and article display.




Figure 3: Summary Window




Figure 4: Article Window




Figure 5: Possible relationships between a query and a database content label.

#### *Content Labels for Query Routing*

Database *content labels* are used by the system to automatically *route* a query to the database or databases where the query can be processed. A content label is a succinct description of the data contained in a database. For example, a content label could state that a database contained all New York Times movie reviews for 1986, or it could state that a database contained all articles by Quindlen in 1989. Content labels and database queries are expressed using predicates.

Content labels are used to route queries and updates. In order to route queries to appropriate databases we use a simple theorem prover. If a query predicate implies a database content label predicate then we know that the query can be processed at the database. Database updates must be sent to all databases whose content label the updates satisfy.

Figure 5 illustrates the different relationships than can hold between a query and a content label. Query Q1 describes a potential result set that is only partially contained in DB1, but is contained entirely in DB2. Thus Q1 is processed at DB2. Query Q2 describes a potential result set that is only partially contained in DB2. In this case the alternative query Q2' is generated and the user is asked for confirmation before Q2' is processed. Query Q3 describes a potential result set that is not contained in DB1 or DB2. Thus Q3 is rejected.

The content label of a personal database is a user's filter list. Because the filter list precisely describes what is in a user's personal computer, the list is used to determine which queries can be processed locally, and which queries must be forwarded to a server.

The content labels of available server computers are kept in a configuration file used by the personal computer's database system. The system will send queries that can not be processed locally to an appropriate server. When a query can not be processed anywhere the system will so inform the user.

Servers in fact are comprised of many independent databases. These independent databases are not visible from outside of the sever, but once a query enters a server a second level of query routing is used to dispatch the query to appropriate internal databases. This organization permits internal databases to be moved between servers and to be moved on and off line as a unit.

#### *Communication Protocols*

Our simplex communication protocol allows us to achieve any desired level of reliability while using an unreliable simplex link [Gifford85a]. Data blocks (such as articles) are fragmented into packets for transmission. Each packet includes a checksum for error detection, and packets are repeated for error correction. Packet repeats are separated in time to minimize the effects of burst errors. We operate with a packet size of 128 bytes with a packet repeat count of 2, allowing us to transmit more than 20 MBytes of data in 24 hours. In our primary service area around Boston the raw bit error rate of our simplex channel is less than  $10^{-6}$ , and after error correction it is less than  $10^{-8}$ .

All the data we broadcast is cryptographically protected against unauthorized use. Each user has a private *key ring* stored in his or her personal computer that describes the data that the user is allowed to receive. Each data block we transmit is encrypted under a nonce key, after the nonce key has been encrypted with a data source specific master key, this nonce key is transmitted along with the block. Our bulk encryption algorithm employs a linear feedback shift register with a non-linear output function. Further details on our broadcast protocols and cryptographic techniques can be found in [Gifford85a].

Our duplex communication protocol provides the advantages of remote procedure call (RPC) while allowing us to make effective use of high-latency and low-bandwidth channels [Gifford88a]. A problem with remote procedure call is that a round-trip delay is incurred for each message that limits the

performance of RPC on moderate performance duplex channels. We have developed the *channel model* to overcome this limitation of RPC. The channel model incorporates remote pipes along with remote procedures. A *pipe* is like a remote procedure, except that calls on a pipe do not return values and do not block a caller. Pipe calls are optimized for maximum throughput, unlike remote procedure calls which are optimized for minimum delay. In order to control the concurrency introduced by pipes we introduce the idea of a *channel group*. All the calls on a channel group made by a single caller are serially processed in the order in which they were made.

## 4 Experimental Results

We designed the experimental test of the Boston Community Information System to answer the fundamental question "*Is the information system technology that we have developed useful?*". In order to address this question, we performed a longitudinal study on the use of the system by 200 people outside of our research group.

During late 1982 and early 1983, we worked out the design of an initial system. Once the design was completed in early 1983 we implemented a prototype. This prototype was operational at a site outside of our research group in April 1984, and a small initial test population of 15 users was established. Based upon the experience with this prototype, the system was largely reimplemented during late 1984 and early 1985 to prepare it for a larger scale user test. The initial digital radio channel was replaced with a more reliable design, and the software was enhanced to provide new services.

In 1985, we formulated a plan for a two year test of the Boston Community Information System at approximately 200 local area homes and businesses. The test started on December 17th, 1985 with Version 6.0, a simplex only version of BCIS. On October 7, 1986, users received Version 8.17 of the system, a polychannel version of BCIS. In Version 8.17, queries which cannot be processed at a user's personal computer are automatically forwarded, via modem, to server computers at MIT for processing.

#### 4.1 Experimental Design

The experimental test of the system was performed on a self-selected population of computer literate volunteers. In order to simplify the test subjects were required to own or have access to an IBM Personal Computer. Participants were recruited via advertisements in the MIT paper *Tech Talk*, a mailing to all MIT computer science professors, announcements at the IBM PC users group of the Boston Computer Society, notices in Boston Computer Society publications, and by asking existing participants to recruit new participants. The professional profile of the participants is: .22 students, .21 professors, .15 hardware or software engineers, .09 administrators, .06 research scientists, .04 physicians, and .03 librarians. The remaining .20 of the population is comprised of professionals who are not computer scientists.

Participants were given the system for free. However, each participant was obligated to complete and return a two page questionnaire which we mailed out each month. With the questionnaire, we sent a one page newsletter and a postage paid return envelope. At the onset of the experiment we told participants that they would have to return the Boston CommInS system which was loaned to them if they failed to complete and return their monthly questionnaires. In order to reinforce this stipulation, each participant signed a legal agreement to this end when they joined the experiment, along with a statement that they would not redistribute information received via the experiment. Every month, telephone calls were made to participants who had not sent in their questionnaires for the prior two months. As the experimental data show, some of these users returned their systems and some began again to mail in the questionnaires.

Monthly questionnaires consisted of short answer questions and an essay topic. The short answer questions were designed to focus on particular content, technology, or overall system issues. The essay topics were designed for evaluations and opinions on a wide variety of qualitative issues. All questionnaire responses are confidential.

We designed the questionnaires to separate participants' reactions to the information provided via the system from their reactions to the technology we have developed. Our experimental results show that people were largely pleased with the information provided via the system. Therefore, we believe that the experimental results on the technology aspects of the system do not reflect unhappiness with the content the system provides.

In summary, our results are likely to more favorable than would be reported by the general population of personal computer users because of the selection method for our experimental population. These people are most likely "early adopters" or users who would be likely to use the system initially. However within this constraint we have tried to perform a fair test of the technology we have developed.

## 4.2 Experimental Conclusions

A summary of our experimental results is presented below. In the summary we have included references to questionnaire dates in square brackets, e.g. [February 1986], and averages of 0 to 10 scale ratings are shown with the maximum value of the scale, e.g. 6.6/10. Our 1986 primary data are available in [Gifford87], and our 1987 primary data are available in [Gifford88b].

### *The system provides a useful service*

Our major conclusion is that a polychannel system provides a valuable service, based upon data from the experimental test. Participants rated the overall system as being valuable (6.6/10 [February 1987]) and said they felt better informed with the system (7.4/10 [February 1987]). Figure 6 is a histogram of how valuable participants rated the system [February 1987] on a 0 to 10 scale. Perhaps more importantly participants reported that they actually use the system an average of 31 minutes a day, during which time they scan 46 article summaries, 21 articles, and carefully read 8 articles [February 1987]. Figure 7 is a histogram of how much time participants reported that they actually use the system [February 1987]. Participants reported that they would pay an average of \$12.42 a month [January 1987] for the system, which further supports our conclusion that the system is valuable.

Reported value and use of the system over time shows continued interest in the system. Figure 8 is a plot of the 75th percentile, 50th percentile, and 25th percentile of reported value of the system over a two year period. Figure 9 is a similar plot that shows how many minutes per day users report they use the system. In Figure 9 it is unclear if the drop in the 25th percentile line represents a drop in interest or an increase in user efficiency.

As one might expect users who report the system is valuable tend to use it more. Figure 10 is a scatter plot of reported system value vs. the reported




Figure 6: Histogram of reported value of system on a scale 0 to 10, February 1987,  $N = 105$ .




Figure 7: Histogram of reported time spent sitting at the keyboard and using the system per day in minutes, February 1987,  $N = 103$ .




Figure 8: 75th, 50th, and 25th percentile of reported system value on a scale of 0 to 10 plotted against month of test.





Figure 9: 75th, 50th, and 25th percentile of reported time spent sitting at the keyboard and using the system per day in minutes plotted against month of test.



number of minutes per day of system use.

It appears that users who run the system more find it of greater value. Figure 11 is a scatter plot of reported minutes per day of system operation *vs.* reported system value. It is also possible that users who find the system more useful tend to run it more.


Participants reported that their use of the system was slightly biased towards personal as opposed to professional use (4.1/10 [October 1986]), and that the system caused them to use their personal computers more frequently (7.0/10 [March 1986]). Users felt that the system has more applications in the business world (6.8/10 [February 1987]) than in the home environment (4.8/10 [February 1987]). Participants reported that .3 of the systems were installed in an office [July 1987].

When looking for specific information participants preferred to use the system compared to a newspaper (6.7/10 [February 1987]) but when reading for general information they slightly preferred a newspaper (4.8/10 [February 1987]). Participants reported that they read newspapers less frequently (4.29/10) and watch television news less frequently (4.5/10 where 0 is less frequently, 5 is the same frequency, and 10 is more frequently [February 1987]).

Comments by users include "I think the system is great. I've gotten so used to it the only time I read a newspaper is when I find it on the subway, and even then, I know most of the stories" [August 1986], "We are quickly getting hooked on our daily information gathering via the air waves" [August 1986], and "... my interest in the system remains undiminished. If anything, I have become a more sophisticated user, more likely to add, change, or remove filter lines temporarily to capture articles of transient interest..." [August 1986]. The most negative comment we received was "I feel the software could be better designed. Also having to dedicate my computer to it is a hassle" [March 1986].

#### *The simplex component of the system is essential*

Participants rated the simplex system as very useful (7.0/10 [April 1987]). The average participant runs the simplex system on his or her personal computer about 8 hours a day, with 25% of the users running it 10 or more hours per day [April 1987]. In their written essays many participants describe the simplex system using terms such as "extremely valuable", "quite useful", and "absolutely vital for the CIS to be successful" [April 1987].



**Figure 11:** Scatter plot of reported time system receiving data per day in minutes (x-axis) vs. reported system value on a 0 to 10 scale, February 1987.

### *The duplex component is important*

The duplex part of the system appears to be an important component that strengthens the system even though participants do not use it frequently. Participants rated the duplex system as somewhat useful (4.3/10 [April 1987]). Thus the duplex system was reported to be less useful than the simplex system. .38 of the participants reported using the duplex system [November 1987]. A contributing factor to this usage may be that .23 of the participants do not own modems [July 1986]. .49 of the users with modems did use the duplex system. The average user reported using the duplex system .8 times a week, with a duplex session lasting an average of 9 minutes [November 1987].

Participants reported that the duplex component may be critical to the commercial success of the system (4.6/10 [May 1987]), and indicated that they would not be happy if the duplex component was eliminated (2.6/10 [May 1987]). In addition, several participants noted the synergism between the simplex and duplex components of the system in their written essays. For example, participants wrote "Broadcasting and remote access are integral parts of the whole system. The interactive (remote access) portion strengthens the passive (broadcasting) role of the system" [May 1987], "I don't use the two-way part much, but I have found it very useful when I did use it. I feel both aspects are important" [May 1987], and "I like the 2-way system and use the system a lot more..." [January 1987].

Interestingly two participants wrote that they found the duplex system much more useful than the simplex system, for example: "Broadcasting for general news is not useful to me. Remote access is useful."

### *Processing of Information*

Participants were enthusiastic about the filtering facilities offered by the system and had an average of 30 lines in their personal filters [March 1986]. Participants wanted the ability to use longer filter lines (7.2/10 [January 1988]) and to turn filter lines on and off (7.5/10 [January 1988]). A common theme in essays was the need to keep refining filters to get the most value from the system. One participant wrote "First and foremost, I've built some skills at writing filters to catch the articles I want" and another wrote "I change, refine, and delete filter lines on a regular basis" [August 1986].

### *System Failings*

The most common complaint about the system was the need to dedicate a personal computer for its operation. Participants reported that they would pay \$90 for a "black box" that would receive articles of interest while their personal computer was turned off [January 1987]. Users also complained that it was difficult to produce appropriate filter lines and that duplicate articles were not adequately filtered out.

In March 1987 we sent a follow-up questionnaire to 90 former participants and we received 42 responses. The reasons most often cited for dropping out of the experiment were lack of time, poor radio reception, inadequate access to an IBM personal computer, and fear that continuous operation would damage a personal computer.

## 5 Conclusions

We have proposed a polychannel architecture for future distributed systems which is intended to combine the scale and cost advantages of simplex communication with the flexibility of duplex communication. Our experimental test of The Boston Community Information System has confirmed our hypothesis that this architecture can be used to provide a new type of useful service. Users reported that they found both the simplex and duplex parts of the system valuable, with most user requests being satisfied using data received via simplex transmission. Users liked the personalized filtering of information which created an active medium.

We expect that polychannel systems will become even more attractive in the future. Within 10 years the processor and memory of a complete IBM-PC class machine will be on a single VLSI part. This part will allow an IBM-PC class computer to be integrated into a standard telephone-like appliance. Emerging storage technologies will make it possible for us to use a polychannel system for creation of personal digital libraries. Digital broadcast to the home and office will be easier to accomplish via future versions of the telephone local loop. In addition, there will be a greater need for computer aided processing of information, since fiber optic technology will allow much more information to be delivered to us.

We expect that the local computational power of polychannel systems will be used to interpret documents that contain computational elements. Appli-

cation programs such as Hypercard make it very easy for non-programmers to create active documents. In the future such active documents will be an important communication medium because they will allow us to represent knowledge with computational elements which are more reactive than is linear text on paper.

In sum, our experimental findings in the context of present technology trends suggest that polychannel systems will offer an important future alternative to traditional fully duplex distributed system designs. As our data show these future systems may provide us with a new kind of personalized mass media. This mass media would offer us the ability to learn more about our world and to communicate more effectively with one another.

## References

[Gifford85a] Gifford, D.K., Lucassen, J.M., and Berlin, S.T., "The Application of Digital Broadcast Communication to Large Scale Information Systems", IEEE Journal on Selected Areas in Communications SAC-3, 3 (May 1985), pp. 457-467.

[Gifford85b] Gifford, D.K., Lucassen, J.M., and Berlin, S.T., "An Architecture for Large Scale Information Systems", Proc. 10th ACM Symposium on Operating System Principles, Orcas Island, Washington, USA, December 1-4, 1985, pp. 161-170.

[Gifford87] Gifford, D.K., *et al.*, "Boston Community Information System 1986 Experimental Test Results", MIT/LCS/TR-397, MIT Laboratory for Computer Science, Cambridge, MA, August 1987.

[Gifford88a] Gifford, D.K., and Glasser, N., "Remote Pipes and Procedures for Efficient Distributed Communication", ACM Trans. on Computer Systems, August 1988.

[Gifford88b] Gifford, D.K., *et al.*, "Boston Community Information System 1987 Experimental Test Results", MIT/LCS/TR-422, MIT Laboratory for Computer Science, Cambridge, MA, August 1988.

[Lindsay84] Lindsay, B.G., *et al.*, "Computation and Communication in R\*: A Distributed Database Manager", ACM Trans. on Computer Systems 2, 1 (February 1984), pp. 24-38.

[Lotus87] "Lotus Signal", Lotus Information Network Corporation, San Mateo, CA, 1987.

[Lucassen86] Lucassen, J., *et al.*, "Boston Community Information System User's Manual (Version 6.0)", MIT/LCS/TR-352, MIT Laboratory for Computer Science, Cambridge, MA, April 1986.

[Segal86] Segal, D., *et al.*, "Boston Community Information System User's Manual (Version 8.17)", MIT/LCS/TR-373, MIT Laboratory for Computer Science, Cambridge, MA, September 1986.

[Tanton79] Tanton, N.E., "Teletex - Evaluation and potential", IEEE Trans. Consum. Electronics, vol. CE-25 (July 1979), pp. 246-250.

**OFFICIAL DISTRIBUTION LIST**

Director Information Processing Techniques Office Defense Advanced Research Projects Agency 1400 Wilson Boulevard Arlington, VA 22209	2 copies
Office of Naval Research 800 North Quincy Street Arlington, VA 22217 Attn: Dr. R. Grafton, Code 433	2 copies
Director, Code 2627 Naval Research Laboratory Washington, DC 20375	6 copies
Defense Technical Information Center Cameron Station Alexandria, VA 22314	12 copies
National Science Foundation Office of Computing Activities 1800 G. Street, N.W. Washington, DC 20550 Attn: Program Director	2 copies
Dr. E.B. Royce, Code 38 Head, Research Department Naval Weapons Center China Lake, CA 93555	1 copy